Challenges with BGPSec

Jan Oesterle, Holger Kinkelin*, Filip Rezabek*
*Chair of Network Architectures and Services, Department of Informatics
Technical University of Munich, Germany
Email: ge25goc@mytum.de, kinkelin@net.in.tum.de, frezabek@net.in.tum.de

***Abstract*—BGP serves as the standard inter-domain routing protocol. It exchanges Network Layer Reachability information between Autonomous Systems and by this ensures connectivity across the Internet. At the time BGP was introduced, there were no security concerns. The inadequate security led to numerous attacks on the Internet, the paper covers. The lack of security resulted in multiple different attempts to fix this issue. One of these attempts is BGPSec. This paper explains this extension to BGP and discusses the degree of security it offers. Because additional security comes with an additional cost, this paper analyzes the deployment issues that exist. In conclusion, it was found that BGPSec is a good start as it solves some existing vulnerabilities. Nevertheless, it is still a work in progress as there are still vulnerabilities and high deployment costs.**

***Index Terms*—border gateway protocol, bgpsec, resource public key infrastructure**

1. Introduction

The modern Internet consists of multiple smaller networks, the so-called Autonomous Systems (AS). AS are administered by a single organization and are reachable by an IP prefix. These networks can, for example, be companies, local internet providers, or universities. To identify individual Autonomous Systems, each of them gets assigned a globally unique number. These numbers are administered by the Internet Assigned Numbers Authority (IANA) and assigned to Regional Internet Registries (RIR), who assign them further.

Due to this distributed nature of the Internet, there is a necessity for routers to exchange information about networks they can reach, allowing them to decide where to forward received packets. This exchange of information is called routing. Routing between AS is called external routing, and BGP is the de-facto standard protocol used for this. Over time it evolved to its current 4th version as described in RFC 1105 [1]

While creating a high standard of interconnectivity, BGP lacks in ensuring security. Over time, this leads to some devastating effects globally due to either accidental misconfiguration or malicious intent. Since then, multiple approaches to add security to the protocol were formulated. BGPSec is one of these approaches and the topic of this paper.

The rest of the paper is structured as follows. In Chapter 2, the routing process of BGP is explained. Chapter 3 analyzes the vulnerabilities of the current state of BGP. Afterward, Chapter 4 follows an introduction to BGPSec, focusing on what it tries to achieve. Chapter 5 compares what vulnerabilities BGPSec solves and what attacks are still possible. In chapter 6, this is accompanied by a discussion of the deployment hurdles BGPSec has to overcome to become the new standard. The paper ends with a conclusion on whether the additional security justifies the effort that has to be taken to deploy BGPSec in Chapter 7.

2. The BGP Routing Process

The BGP belongs to the family of path-vector protocols. In path-vector protocols, the most important exchanged routing information is a destination, and path packets have to traverse to reach this destination. Destinations come as an IP prefix, and paths come in the form of a list of AS numbers. The exchanged information is called Network Layer Reachability Information (NLRI).

For two routers to be able to exchange NLRI, they first have to establish a direct connection. This connection is built upon a Transmission Control Protocol (TCP) connection and called BGP peer relationship. To establish this peer relationship, the two peers exchange `OPEN` messages to negotiate parameters of a peer relationship. Such parameters are, for example, capabilities like the use of BGPSec or a maximum time interval the connection will be kept open in case they do not exchange messages. This time interval is called hold-timer and is used to evaluate whether a peer relationship is still active. If the peers do not exchange messages for one full hold-timer, the connection between the two peers is closed. This closing leads to them dismissing all routing information they gained from this connection. To prevent this, peers regularly exchange `KEEPALIVE` messages to reset the hold timer. `UPDATE` messages carry the actual NRLI. The last class of messages specified by BGP is `NOTIFICATION` messages. Peers use these messages to inform other peers about possible errors such as malformed packets.

Figure 1 shows an exemplary routing process. AS1 announces in messages 1) and 2) the prefix 192.0.20./24 to both its peers AS2 and AS3. In the red path, after receiving message 1), AS2 prepends its own AS number to the path and sends message 3) to AS5 with the updated path attribute. In the green path, both AS3 and AS4 prepend their AS number, as one can see in messages 4) and 5). This routing example results in AS5 receiving the two `UPDATE` messages 3) and 5), announcing the same prefix. Because the prefix of both of the messages is identical, AS5 can choose what path to prefer. The router could base the decision on the length of the path leading

doi: 10.2313/NET-2022-01-1_02

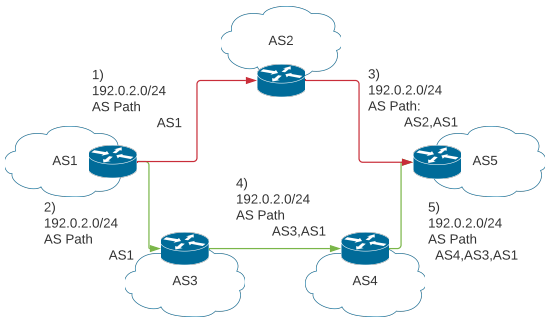

Figure 1: BGP routing process

to AS5 preferring the red path. Another deciding factor could be the economic relationship of the peers. Consider the example of AS2 being a provider and AS4 being a consumer of AS5. AS5 then would have to pay AS2 for traffic but get paid from AS4 for traffic. This monetary difference may lead to a preference for the green path despite it being longer. This decision process is called policy-based routing and can be configured by the router administrator.

3. BGP - Security Concerns

BGP version 4 (RFC 4271) [2], addresses connectivity and scalability demands but makes no considerations towards security. This lack of security makes it easy for accidental or malicious misconfiguration that can devastate the Internet as a whole. As BGP uses TCP as the underlying protocol, TCP's known weaknesses can serve as an additional attack vector. Possible attacks on TCP include attacks such as eavesdropping, insert forged BGP `UPDATE` messages, and Denial of Service attacks. [3]

3.1. Prefix Hijacking

Prefix hijacking is a common attack type in BGP. Its goal is to hijack traffic headed to a specific destination by announcing a more specific prefix to the destination's one. This attack exploits the mechanism of more specific prefix matching.

More specific prefix matching is a standard in routing and used to choose a fitting entry in a routing table in case an IP address matches more than one entry. Consider the example of a router with two entries 1) Destination: 123.4.5.0/24 Next Hop: AS3 and 2) Destination: 123.4.0.0/16 Next Hop: AS5. In the case this router receives a packet with the destination IP 123.4.5.6, it has to choose to what AS it forwards the packet to as both entries fit this IP address. In this case, the router executes more specific prefix matching by preferring the "longest" prefix. In this example, the packet would be forwarded to AS3.

Prefix hijacking makes use of routers' ability to announce arbitrary prefixes and more specific prefix matching. This allows a router to hijack traffic bound to a prefix by announcing a more specific version of it.

Prefix hijacking can be divided into two categories. 1) Black Hole attacks and 2) Interception attacks. The difference between them both is the way they handle the hijacked traffic. In Black Hole attacks, the traffic is attracted and then dropped. Instead of dropping the packets, Interception attacks forward them to the original destination creating a Man in the Middle (MitM) attack enabling the attacker to read and alter packets.

One prominent example of a Black Hole attack is the Pakistan Youtube hijack. This also serves as a good example that even accidental misconfiguration can cause great harm. In 2008, Pakistan made plans to block Youtube country-wide [4]. The Pakistani government instructed Pakistan Telekom to realize this block. They attempted to announce a more specific prefix than the one Youtube announced. and by this, attracting all the traffic originally bound to Youtube. A simplified structure of this attack can be seen in Figure 4.

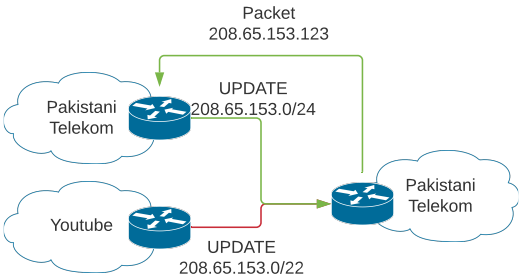

Figure 2: Pakistani Youtube Hijack

Pakistan Telekom announced 208.65.154.0/24. Because of the global propagation of prefixes and /24 being more specific than /22, they got preferred by most existing BPG routers. The green path indicates this. This brought Youtube eventually down for about two hours. The hijack was solved by Youtube announcing even more specific prefixes, effectively hijacking their traffic back.

3.2. Impact on the Internet

As seen in the Youtube hijack incident, a single announcement can greatly impact the Internet. As the Internet traffic steadily grows, so does the amount of sensitive data on the Internet. Recent BGP hijacks show that primary goals were companies that hold vast user data, such as Amazon, Facebook, Google, and Banks. [5] This makes securing BGP a significant concern.

4. BGPSec - an Extension to BGP

Efforts to address BGPs vulnerabilities led to a multitude of different approaches over time. One proposal was to introduce path validation to the protocol. The BGP Security Extension in RFC 8205 [6] formulates this proposal.

4.1. Goals of BGPSec

The introduction of path validation and origin validation intends to harden BGP to achieve byzantine robustness. Byzantine robustness is described as in case

of malicious or faulty behavior of hosts, the other hosts should 1) receive the same message that was sent by the original host 2) reach a decision on a message's contents within a finite time period 3) this decision should be the same among all these hosts [3].

4.2. Path and Origin Validation

Path validation is a mechanism that allows routers to validate the path information contained in `UPDATE` messages. The validation checks whether the announced path matches the actual path packets will take. Origin validation asserts whether the announcing AS owns the prefix contained in the `UPDATE` message. For routers to execute these validations, an additional infrastructure is needed that holds information about AS numbers and prefix owners. A possible implementation of such an infrastructure is called Ressource Public Key Infrastructure.

4.3. Resource Public Key Infrastructures

The Resource Public Key Infrastructure (RPKI) is used to issue and distribute certificates that link resources to resource holders. Such resources can be IP prefixes and AS numbers. [7]. These certificates are then published and made available for the public on dedicated repository servers. These certificates can be queried for matching AS numbers, IP prefixes, and Subject Key Identifiers (SKI). These are identifiers used in case an AS number corresponds to multiple certificates.

4.4. Certificate Issuing Process

The certificate issuing process is hierarchical and in accordance with the allocation of IP address space. Consider the following example: IANA allocates the address space 123.0.2.0/24 and the AS number 20 to the Regional Internet Registry RIPE NICC. RIPE NICC, in turn, allocates the AS number and address space to a university network.

IANA then would assign a certificate to RIPE NICC holding the authority to use the AS number 20 and a certificate holding the authority to announce IP prefixes in 123.0.2.0/24. As these certificates authorize an entity to announce particular prefixes, they are called Route Origin Authentication (ROA). RIPE NICC subsequently issues another set of certificates to the university and publishes the certificates in a publicly accessible repository server.

4.5. Exemplary BGPSec Routing Process

In BGPSec, the `AS PATH` attribute gets replaced with the `BGPSec PATH` attribute to hold the additional information in the form of a signature block. Each signature in that block corresponds to an AS number in the path. So the longer the path gets, the more signatures such a block will contain. The way signatures are created is based on whether the router announces a prefix or propagates routing information. A prefix-announcing router create the signature based on the announced prefix, their own AS number, and the AS number they forward the `UPDATE` message to. On the other hand, a router that propagates a received `UPDATE` message uses the previous signature instead of a prefix to create a new signature. Each of these signatures is accompanied by a SKI.

Figure 3 shows an exemplary routing process. AS1 announces its prefix 192.0.2.0/24. It begins by prepending its AS number to the `Secure Path` and then create a signature block corresponding to its AS number. When AS2 receives the `UPDATE` message sent by AS1, it validates all signatures in the signature block and then appends its own AS number to the `BGPSec Path` and a new signature to the signature block. After this, it forwards the `UPDATE` message to AS3. The router at AS3 then again validates the information. As now two signatures are contained in the signature block, the router at AS3 has to validate two signatures. It begins with the most recent one, in our example sig2. To validate it, it queries the certificate that matches SKI2 and AS2. If a matching certificate is found, they use the public key to validate the signature cryptographically. If this validation fails or no certificate was found, the `UPDATE` message will be deemed invalid. Then, the router validates the next signature the same way. After validating the last signature, the router can query the ROA corresponding to the contained prefix.

Based on the validated signatures, the router can ensure that each AS number in the path belongs to the router that created the signature. Furthermore, the router can ensure that the next hop in the signature corresponds to the next AS number in the path. By this, path validation is achieved. Moreover, by querying the ROA, the router can ensure that the origin AS is allowed to announce the prefix. With this, Origin Validation is ensured.

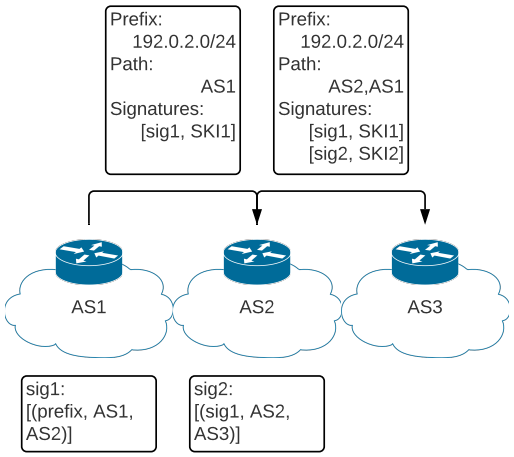

Figure 3: BGPSec Routing Process

5. Analysis of the Effectiveness of BGBPsec

As mentioned above, BGPSec only ensures valid origins and that paths are genuine. While preventing attacks to some degree, there are still vulnerabilities that have to be addressed.

5.1. What It Prevents

BGPSec covers most accidental misconfiguration and unsophisticated attacks. For example, the Youtube hijack

doi: 10.2313/NET-2022-01-1_02

from chapter 3 would be prevented. The first BGP router that receives the UPDATE message would try to validate the signature. This validation attempt would fail because Pakistani Telekom did not use Youtube's private key as it is, at least in theory, not in their possession.

5.2. What It Does Not Prevent

One major issue with BGPSec as it is at the moment is, that traffic hijacking is still possible. A good example is the wormhole attack. The basic idea of that attack is to create a shorter path than the current one and by this redirecting the traffic. Figure 4 shows such an attack.

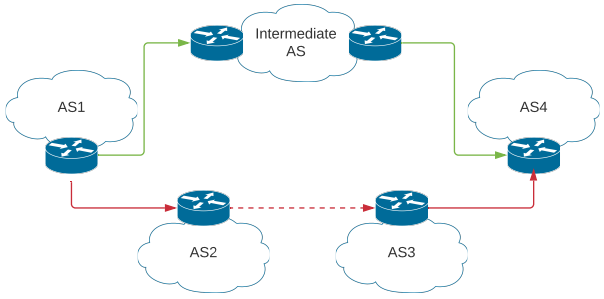

Figure 4: Wormhole Attack

To conduct a wormhole attack, an attacker needs to control two BGP speakers´. These have to be in a peer relationship with the endpoints of the traffic the attacker wants to hijack. In this example, the attacker is in control of speakers at AS2 and AS4. Before the attack, the traffic between AS1 and AS4 is forwarded via the green path. The attacker now creates a tunneled peer relationship between AS2 and AS3 indicated by the red dashed line. This creates a path of length three between AS1 and AS4. In about 86 percent of cases [8], path lengths are longer than 3. This leads to the red path being preferred due to shorter path length and leads to a MitM attack. Because the attacker does not announce a prefix and does not forge the BGPSec PATH attribute, this attack is not covered by BGPSec.

6. Discussion of Deployment Issues

Additional security comes with an additional cost in terms of storage and processing power requirements. This cost and the still obvious flaws BGPSec has played a major role in the hurdles it has to overcome in terms of deployment.

6.1. A Technical View

The introduction of signatures and certificates to the routing process is crucial for the additional security BGPSec offers. Each time a BGP router receives an UPDATE message, it validates all signatures. Additionally, each time a router propagates an UPDATE message, it has to create a signature and append it to the signature block in the BGPSec Path. This leads to an increasing number of validations the more extended the path gets, and makes it on average 70 times slower than regular BGP [8]. The number of required UPDATE message itself also increases. This is because BGPSec Path attributes can only contain one single prefix while regular AS PATH attributes can contain multiple ones [6]

This serves as a significant hurdle because, according to RFC 7747 [7], convergence is a major factor in the reliability of BGP. Convergence means that all routers have the same information about the network topology. Due to the continuous change of this topology, the propagation of this change should happen fast. Because of BGPSec's longer processing time, convergence is slower than the convergence without it.

The upside of BGPSec is that as an extension, it can work in parallel with regular BGP. As stated in the RFC, the BGPSec Path attribute is an optional attribute that replaces the AS path attribute. The decision of what attribute to use is negotiated between peers using the OPEN messages. In the case that a BGP speaker wants to propagate a prefix is received from a peer connected using BGPSec, the BGPSec path attribute will get stripped of its additional information and then propagated as AS path to the peers that do not use BGPSec. This allows for a gradual deployment because BGP and BGPSec routers can coexist, and communication between them does not affect the routing process. in the scope of insecure routing.

6.2. A Management View

As stated in RFC 4271 [2], the management of certificates and origin/prefix pairs are handled by two distinct RPKIs. Because BGP is not under a single authority, collecting complete data sets and keeping them up to date is a significant deployment hurdle BGPSec still has to face. At the moment, there is already an RPKI in place for origin authentication. According to a report on RPKI [9], about 27 percent of prefix announcements are valid, 0.5 invalids, and 72,4 unknown, meaning that the RPKI has no information about the pairing of prefixes to AS numbers. This shows that there were some efforts to implement it, but wide-scale deployment has yet to be achieved. Additionally, there is no incentive to provide such data for a single ISP because they get no direct value out of this. An approach to change this may be the deployment beginning with more prominent parts of the Internet and discrimination of the ones that did not implement it by preferring connections through paths using BGPSec.

7. Conclusion

This paper presented BGP as the de-facto standard routing process to exchange routing information between AS. As BGP has no built-in security, it is vulnerable to attacks such as the famous Pakistan Youtube incident. BGPSec is a proposed extension to BGP that adds path and origin validation to the routing process by using RPKI. However, although it prevents some attacks and misconfiguration from happening, there are still significant flaws that allow for attacks like wormhole attacks. These vulnerabilities show the state of the protocol as a work in progress. Contributing to that are the still prevalent issues it faces in terms of deployment.

The additional security BGPSec offers comes with a price. The creation of signatures and the validation process takes more time and needs additional space. Furthermore, it is hard to collect and manage the necessary data, as there is no single authority that manages AS numbers and prefixes. With BGP being an old protocol, these problems follow the problems other protocols of that era face, like DNS. With approaches to improve BGPSec existing but not yet included in the current RFC, more research is still necessary to find a fitting solution to the current problems BGP is facing.

References

[1] K. Lougheed and J. Rekhter, "Border gateway protocol (bgp)," Internet Requests for Comments, RFC Editor, RFC 1105, June 1989, http://www.rfc-editor.org/rfc/rfc1105.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc1105.txt

[2] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," Internet Requests for Comments, RFC Editor, RFC 4271, January 2006. [Online]. Available: http://www.rfc-editor.org/rfc/rfc4271.txt

[3] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A survey of bgp security issues and solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010.

[4] (2008) Youtube hijacking: A ripe ncc ris case study. [Online]. Available: https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

[5] J. Sherman. (2020) The politics of internet security: Private industry and the future of the web. [Online]. Available: https://www.atlanticcouncil.org/in-depth-research-reports/report/the-politics-of-internet-security-private-industry-and-the-future-of-the-web/

[6] M. Lepinski and K. Sriram, "BGPsec Protocol Specification," Internet Requests for Comments, RFC Editor, RFC 8205, September 2017. [Online]. Available: http://www.rfc-editor.org/rfc/rfc8205.txt

[7] M. Lepinski and S. Kent, "An infrastructure to support secure internet routing," Internet Requests for Comments, RFC Editor, RFC 6480, February 2012, http://www.rfc-editor.org/rfc/rfc6480.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc6480.txt

[8] K. Kim and Y. Kim, "Comparative analysis on the signature algorithms to validate as paths in bgpsec," in *2015 IEEE/ACIS 14th International Conference on Computer and Information Science (ICIS)*, 2015, pp. 53–58.

[9] (2020) Global prefix/origin validation using rpki. [Online]. Available: https://rpki-monitor.antd.nist.gov/

doi: 10.2313/NET-2022-01-1_02